arXiv:1704.04509v1 [cs.DS] 14 Apr 2017

The Entropy of Backwards Analysis

Mathias Bæk Tejs Knudsen*†and Mikkel Thorup*

University of Copenhagen
mathias@tejs.dk, mikkel2thorup@gmail.com

Abstract

Backwards analysis, first popularized by Seidel, is often the simplest most elegant way of analyzing a randomized algorithm. It applies to incremental algorithms where elements are added incrementally, following some random permutation, e.g., incremental Delauney triangulation of a pointset, where points are added one by one, and where we always maintain the Delauney triangulation of the points added thus far. For backwards analysis, we think of the permutation as generated backwards, implying that the ith point in the permutation is picked uniformly at random from the i points not picked yet in the backwards direction. Backwards analysis has also been applied elegantly by Chan to the randomized linear time minimum spanning tree algorithm of Karger, Klein, and Tarjan.

The question considered in this paper is how much randomness we need in order to trust the expected bounds obtained using backwards analysis, exactly and approximately. For the exact case, it turns out that a random permutation works if and only if it is minwise, that is, for any given subset, each element has the same chance of being first. Minwise permutations are known to have $\Theta(n)$ entropy, and this is then also what we need for exact backwards analysis.

However, when it comes to approximation, the two concepts diverge dramatically. To get backwards analysis to hold within a factor α, the random permutation needs entropy $\Omega(n/\alpha)$. This contrasts with minwise permutations, where it is known that a $1+\varepsilon$ approximation only needs $\Theta(\log(n/\varepsilon))$ entropy. Our negative result for backwards analysis essentially shows that it is as abstract as any analysis based on full randomness.

1 Introduction

Randomization is a powerful tool in the construction of algorithms and data structures, yielding algorithms that are often much simpler and more efficient than their deterministic counterparts. However, much of the analysis relies on the assumption of full access to randomness, e.g., a sequence of n i.i.d. random variables, a hash function mapping keys from a large universe into independent hash values, or as in this paper, a uniformly random permutation $\pi \in S_n$. However, often n is large, and then it is unrealistic to get access to just n independent random bits. A lot of work has been devoted to removing the unrealistic assumption of full randomness, replacing it with

*Research partly supported by Advanced Grant DFF-0602-02499B from the Danish Council for Independent Research under the Sapere Aude research career programme.

†Research partly supported by the FNU project AlgoDisc - Discrete Mathematics, Algorithms, and Data Structures.

weaker notions of pseudo-randomness where the above large objects are not uniformly random, but rather generated based on only few fully random bits. We also want the generation to be efficient, e.g., for a pseudo-random hash function, we want to quickly compute the hash value of any given key. While the pseudo-random objects are not uniformly random, we do want them to satisfy probabilistic properties that can be proved to suffice in many algorithmic contexts. A classic example is Wegman and Carter's *k-independent hash functions* [14] with the property that any k given keys get independent random hash values. We can get such a hash function over the prime field $\mathbb{Z}_p$ using a random degree $k-1$ polynomial, instantiated with k independent random coefficients from $\mathbb{Z}_p$. The important point here is that if we have an application of hash functions where the analysis only assumes uses k-independence, then we know that it can be implemented using any k-independent hash function. Another example more relevant to this paper is minwise permuations of Broder et al. [1]. A random permutation $\pi \in S_n$ is *(approximately) minwise* if for any given subset $X \subseteq [n] = \{1, \ldots, n\}$, each $x \in X$ has (approximately) the same chance of being first from X in π. If a minwise permutation is used to pick the pivots in Quicksort [5], then we know that Quicksort will have exactly the same expected number of comparisons as if a fully random permutation was used—two keys x and y get compared if and only if one of them is picked before any key between them. Exact minwise permutations are unrealistic to implement in the sense that they have entropy at least $\Theta(n)$ [1], which is only slightly less than a fully random permutation that has entropy $\ln n! \approx n \ln n$. However, if we are satisfied that each $x \in X$ is first with probability at most $(1+\varepsilon)/|X|$, then the required entropy is only $\Theta(\log(n/\varepsilon))$ [1], and then the expected number of comparisons in Quicksort is only increased by a factor $1+\varepsilon$.

Randomized algorithms have now been studied for more than 50 years [5, 11]. We have a vast experience for how randomized algorithms can be understood and analyzed. Probabilistic properties have emerged that have proven useful in many different contexts, and we want to know if these properties can be supported with realistic implementations. Generally, we want the simplest most elegant analysis of a randomized algorithm, but if it relies on unrealistic probabilistic properties, then we have to also look for a possibly more complicated analysis relying on more realistic properties, or worse, give up the algorithm because we do not know how to implement the randomness.

1.1 Backwards analysis

In this paper we focus on, so-called, backwards analysis, which was first popularized by Seidel [12]. Using backwards analysis he obtained an extremely simple and elegant analysis for many randomized algorithms assuming full randomness. The technique is so appealing that it is now taught in text books [3, 8].

Backwards analysis applies to incremental algorithms where elements from some set $X = \{X_1, ..., X_n\}$ are added one by one in a random order, starting from the empty set. Let π be the corresponding permutation over $[n] = \{1, \ldots, n\}$, that is, $X_{\pi(i)}$ is the ith element added to the set. We note that backwards analysis has also been applied to the case where elements are added in batches [2] rather than one by one. We shall return to this in Section 1.4, and for now just focus on permutations.

We also have a cost function that with any $Y \subseteq X$ and $x \in Y$ associates a cost $c(x, Y)$ of adding x to $Y \backslash \{x\}$. A classic example from [12] is incremental Delauney triangulation of a point set where $c(x, Y)$ is the degree of x in the Delauney triangulation of Y. For a given permutation π, the total cost is $\sum_{i=1}^{n} c(X_{\pi(i)}, X_{\pi([i])})$. Here $[i] = \{1, \ldots, i\}$, $\pi(I) = \{\pi(j) \mid j \in I\}$, and $X_J = \{X_j \mid j \in J\}$, so $X_{\pi([i])} = \{X_{\pi(1)}, \ldots, X_{\pi(i)}\}$. Now, what is the expected total cost if the permutation π is fully-

random?

In backwards analysis, we think of the permutation π as generated *backwards* starting from $\pi(n)$. Then $\pi(i)$ *is uniformly distributed in* $\pi([i]) = \{\pi(1), \ldots, \pi(i)\} = \{1, \ldots, n\} \backslash \{\pi(i+1), \ldots, \pi(n)\}$. The important point here is that when $\pi(i)$ is generated, we only know the set $\pi([i])$ via the values of $\pi(i+1), \ldots, \pi(n)$—we do not now the individual values of $\pi(1), \ldots, \pi(i)$. For Delauney triangulation this means that the ith point $X_{\pi(i)}$ is uniformly distributed amongst the points in the given set $X_{\pi([i])}$. The expected cost is therefore the average degree of a point in the (assumed unique) Delauney triangulation of $X_{\pi([i])}$. Since the Delauney triangulation is a planar graph, by Euler's formula, the average degree is less than 6. The expected cost of adding all n points is thus at most $6n$.

In a more general use of backwards analysis, suppose for $i = 1, \ldots, n$, we have determined c_i such that for any given subset $Y \subseteq X$ of size i, if x is uniformly distributed in Y, then $\mathrm{E}[c(x, Y)] \leqslant c_i$. Then, with a fully random permutation, the expected total cost is bounded by $\sum_{i=1}^{n} c_i$. For Delauney triangulation, we had $c_i = 6$ for all i. However, there are also applications where the c_i are different, e.g., in the analysis of quicksort from [12], we get that the expected number of comparisons with the ith pivot is bounded by $c_i \leqslant 2n/i$, hence that the expected total number of comparisons is bounded by $2nH_n$.

Our question in this paper is how much randomness we need to create a distribution on permutations so that the expected bounds derived from backwards analysis apply exactly or approximately. Based on the many positive findings surveyed in [13], we were originally hoping for a good solution using $O(n^{\varepsilon})$ random bits for some small $\varepsilon > 0$.

We will now define more precisely the probabilistic properties needed for backwards analysis. To simplify notation, we identify our set X of elements with $[n] = \{1, \ldots, n\}$, that is, $x_i = i$ for $i = 1, \ldots, n$. A *cost function over* $[n]$ assigns a cost $c(x, Y)$ to any subset $Y \subseteq [n]$ and $x \in Y$. The total cost of a permutation $\pi \in S_n$ is then

$$c(\pi) = \sum_{i=1}^{n} c(\pi(i), \pi([i])).$$

For our most positive results, we would like our distribution to be backwards α-uniform, defined as follows.

Definition 1. A distribution $\mathcal{D}$ over S_n is *backwards α-uniform* if for any set $Y \subseteq [n]$ with positive support, that is, $\Pr_{\pi \sim D}[\pi([|Y|]) = Y] > 0$, and any $x \in Y$, we have that

$$\Pr_{\pi \sim D}[\pi(|Y|) = x \mid \pi([|Y|]) = Y] \leqslant \alpha/|Y|.$$

If $\mathcal{D}$ is backwards 1-uniform, we will also say that it is *exact backwards uniform*.

Let c be *any* cost function associating a cost $c(x, Y)$ with any $Y \subseteq [n]$ and $x \in Y$. For every $i \in [n]$, let c_i be such that for every size i subset $Y \subseteq [n]$, we have $\mathrm{E}_{x \sim \mathcal{U}(Y)}[c(x, Y)] \leqslant c_i$. If π is drawn from an α-uniform distribution $\mathcal{D}$, then it is easily seen that the total expected cost is $\mathrm{E}_{\pi \sim D}[c(\pi)] \leqslant \alpha \sum_{i=1}^{n} c_i$, which is α times bigger than the bound $\sum_{i=1}^{n} c_i$ we get assuming full randomness like in backwards analysis. Demanding α-uniformity is, however, very conservative. All we really need from the distribution $\mathcal{D}$ is that $\mathrm{E}_{\pi \sim D}[c(\pi)] \leqslant \alpha \sum_{i=1}^{n} c_i$ for any given c and c_is satisfying the conditions. We are going to prove lower bounds for this condition even for the restricted case where all c_i are the same and normalized to 1. We refer to this as being *backwards α-efficient* as defined below.

Definition 2. A distribution $\mathcal{D}$ over S_n is *backwards α-efficient* if the following condition is satisfied. Let c be *any* cost function that with any $Y \subseteq [n]$ and $x \in Y$ associates a cost $c(x, Y)$ such that for every subset $Y \subseteq [n]$, we have $\mathrm{E}_{x \sim \mathcal{U}(Y)}[c(x, Y)] \leqslant 1$. Then

$$\mathrm{E}_{\pi \sim D}[c(\pi)] \leqslant \alpha n.$$

If $\mathcal{D}$ is backwards 1-efficient, we will also say that it is *exact backwards efficient*.

If the distribution $\mathcal{D}$ is backwards α-uniform, it is clearly also backwards α-efficient. Note in the definition of α-efficient that the distribution $\mathcal{D}$ over S_n is given first, and that it has to work for any later cost function. The vision is that we want a "standard library" implementation of the distribution that gives the right expectation for any future cost function. Unfortunately, our main result is negative: as described below, if $\mathcal{D}$ is α-efficient, then we need $\Omega(n/\alpha)$ random bits to define a random permutation π from $\mathcal{D}$.

1.2 Results

We now list our results in more detail. First, concerning exact backwards analysis, we show that a backwards 1-uniform distribution can be implemented as a uniform distribution over a family of $\exp(O(n))$ permutations. We have a matching lower bound showing that even if we allow non-uniform distributions and even if we are satisfied with 1-efficiency, then the family has to be of size $\exp(\Omega(n))$.

Our more interesting results are for α-approximations for $\alpha > 1$. On the positive side we present a backwards α-uniform distribution which is a uniform distribution over a family of size $\exp(O(n(\log \alpha)^2/\alpha))$. On the negative side, we show that even if we allow non-uniform distributions and even if we are satisfied with α-efficiency, then the family has to be of size $\exp(\Omega(n/\alpha))$. More precisely, we show that if a random permutation is α-efficient, then its entropy is $\Omega(n/\alpha)$.

The entropy lower bound on the family size needed for backwards α-efficient permutations is our main result. For a given distribution $\mathcal{D}$, we define a simple adversarial cost function $c^{\mathcal{D}}$ as follows. Consider a set $Y \subseteq [n]$. If $Y \subseteq [n]$ does not have positive support in $\mathcal{D}$, that is, $\Pr_{\pi \sim D}[\pi([|Y|]) = Y] = 0$, then we just set $c(y, Y) = 1/|Y|$ for all $y \in Y$. If Y has positive support, we let $x \in Y$ maximize the chance that it is last in Y, that is, x maximizes $\Pr_{\pi \sim D}[\pi(|Y|) = x \mid \pi([|Y|]) = Y]$. We then set $c^{\mathcal{D}}(x, Y) = |Y|$ and $c^{\mathcal{D}}(y, Y) = 0$ for $y \in Y \backslash \{x\}$. Then $\mathrm{E}_{x \sim \mathcal{U}(Y)}[c^{D}(x, Y)] \leqslant 1$ for all Y as required. If $\mathcal{D}$ is backwards α-efficient, then the expected total cost should be $\mathrm{E}_{\pi \sim D}[c^{\mathcal{D}}(\pi)] \leqslant \alpha n$. However, here we prove that $\mathrm{E}_{\pi \sim D}[c^{\mathcal{D}}(\pi)] \leqslant \alpha n$ implies that the entropy of $\mathcal{D}$ is $\Omega(n/\alpha)$.

1.3 Comparison with minwise permutation

Our results are contrasted with corresponding results for minwise permutations [1]. An *ε-minwise* distribution $\mathcal{D}$ over permutations from S_n is one such that for any given $Y \in [n]$ and any $x \in Y$, we have

$$\Pr_{\pi \sim D}[\pi(x) = \min \pi(Y)] = (1 \pm \varepsilon)/|Y|.$$

If $\varepsilon = 0$, we say $\mathcal{D}$ is exact minwise. In our condition of backwards α-uniform permutations, we only have to consider sets Y with positive support, and then we want

$$\Pr_{\pi \sim D}[\pi(x) = \max \pi(Y) \mid \pi([|Y|]) = Y] \leqslant \alpha/|Y|.$$

Above, the difference between using minimum and maximum is inconsequential. The important difference is the conditioning $\pi([|Y|]) = Y$.

In this paper, we show that for a random permutation $\mathcal{D}$, the following properties are equivalent:

- $\mathcal{D}$ is exact minwise.
- $\mathcal{D}$ is exact backwards uniform.
- $\mathcal{D}$ is exact backwards efficient.

From [1] we know that the family size needed for an exact minwise distribution is $\exp(\Theta(n))$, and hence this is the size we need for exact backwards analysis.

However, when it comes to approximation, it turns out that there is a huge difference between minwise and backwards permutations. For ε-minwise distribution, we have very positive results. In [1] it is proved that for ε-minwise permutations, it suffices with a uniform distribution on a family of size $O(n^2/\varepsilon^2)$. In fact, more constructively, Indyk [6] has show that it suffices to use $O(\log(1/\varepsilon))$-independence. These positive results for ε-minwise stand in sharp contrast to our lower bound of $\exp(\Omega(n/\alpha))$ on the family size needed for a distribution over S_n to be backwards α-efficient.

Viewed in terms of random bits, for constant approximations, the above results say that we need only need a logarithmic bits for approximate minwise permutations whereas we for backwards analysis need a linear number of bits.

1.4 Chan's backwards analysis of the Karger-Klein-Tarjan's MST

Chan [2] has made an interesting application of backwards analysis for the case where elements are not added one by one as in, but rather in a batches. The point is that large batches may be handled more efficiently than elements processed one element at the time. The concrete case is the randomized minimum spanning tree algorithm Karger, Klein, and Tarjan [7] which runs in linear expected time. We assume that all edge weights are unique.

The important step in the algorithm is that we have a graph $G = (V, E)$ from which we sample a random subset $S \subseteq E$ of the edges. Recursively, we compute a minimum spanning forest F from of (V, S). Finally, we remove all edges (v, w) from G that are *F-heavy* in the sense v and w are connected in F by a path with no edge heavier than (v, w). In particular, this includes all edges from $S \backslash F$. The question is, how many edges do we expect to remain in G?

Let us say that we start with n nodes and m edges in G. As in [2], we simplify the discussion a bit by saying that S is sampled to have a specific size pm, and we assume that S is sampled uniformly. We claim that the expected number of edges remaining in G after the removal of F-heavy edges is less than n/p.

The proof is using the idea from backwards analysis. Consider adding uniformly random edge $e \in E \backslash S$ to S. Then e is uniformly distributed in $S' = S \cup \{e\}$ just like $\pi(i)$ was uniformly distributed in $\pi([i])$.

We now make the combinatorial observation that e is F-light if and only if e is in the minimum spanning forest F' of (V, S'). There are at most $n - 1$ edges in F, and e is uniform in S', so we conclude that e is F-light with probability $(n - 1)/(pm + 1) < n/(pm)$. Thus we expect $(m - pm)n/(pm) = n/p - n$ edges to remain from $E \backslash S$. Including the at most $n - 1$ edges from F. It follows that the total expected number of edges remaining is less than n/p, as desired.

The question now is if we can support this kind of backwards analysis based generating a set S of some prescribed size pm, followed by an element e. As in the previous backwards analysis, we

have a cost function $c(S', e)$ of e being added last to $S' = S \cup \{e\}$. The cost is 1 if e is light. We argued for any set S' that $\mathrm{E}_{e\sim\mathcal{U}(S')}[c(e, S')] < n/(pm)$, and based on that, we which to conclude that $\mathrm{E}[\sum_{e\in E\setminus S} c(e, S \cup \{e\})] \leqslant |E\setminus S|n/(pm)$.

Stepping back, normalizing, and allowing for an approximation factor α, what we would like in general for this kind of analysis is a distribution $\mathcal{D}$ over subsets S of E of size pm such that if c be *any* cost function such that for any set E' of size $pm + 1$, we have $\mathrm{E}_{e\sim\mathcal{U}(S')}[c(e, S')] \leqslant 1$, then we conclude that

$$\mathrm{E}_{S\sim\mathcal{D}}[\sum_{e\in E\setminus S} c(e, S \cup \{e\})] \leqslant \alpha|E\setminus S|.$$

Unfortunately, we will show that if p is a constant from $(0, 1)$, then this implies that $\mathcal{D}$ has entropy $\Omega(m/\alpha \log(1 + \alpha))$ unless the approximation factor is very large, i.e. $\alpha = \Omega(m)$.

Chan's MST analysis is an example illustrating how backwards analysis can be applied when elements are added in random batches, and where random elements from one batch are only considered against elements in from preceding batches. In the MST, we have only two batches: (1) the initial samples set S, and (2) the remaining edges $E\setminus S$. This is the extreme opposite of the backwards permutations, where elements were added randomly in batches of size 1. Our entropy lower bounds for both of these extreme cases suggests a general hardness.

1.5 Reducing the randomness

Getting a linear number of random bits in order to trust the bounds from backwards analysis within a constant factor is often prohibitive. It is almost as bad as using $\Theta(n \log n)$ bits to get a fully-random permutation.

For a concrete algorithm like the incremental Delauney triangulation, one can still hope for a different analysis based on weaker assumptions on the distributions of the random permutation π. Indeed, Mulmuley [9] has shown that the bounds mentioned by Seidel [12] can be achieved within a constant factor using only $O(\log n)$ bits. Mulmuley's idea is very simple. For the Delauney triangulation he points out that a triangle uvw will appear if and only if there is no point z inside the triangle appearing before the last of u, v, and w in the permutation, and the time used is determined by the total number of triangles appearing. Mulmuley then showed that generating the permutation using an 11-independent function guarantees that the probability that any given triangle appears is at most a constant fraction larger than it would be, had the permutation been generated using full randomness. It follows that Delauney triangulation has asymptotically the same complexity with 11-independence as with full randomness, and full randomness was analysized with backwards analysis. However, Mulmuley's idea does not apply to the randomized MST algorithm from Section 1.4. The point is that the obstruction for $e = (v, w)$ being F-light requires a whole path of light edges from v to w to be sampled for S. Indeed it is an open problem if the sample S from the randomized MST algorithm can be generated with $o(n)$ bits. Pettie et al. [10] have shown that the MST algorithm can be changed to one that needs only $O(\log n)$ bits, but it would have been much more attractive if an efficient implementation followed directly from the backwards analysis.

If we know in advance that we are only going to study a fixed limited family of cost functions $\mathcal{C}_n$, then we can always use the standard trick of first generating a limited number of fully-random permutations $\Pi_n \subseteq S_n$, that we are free to use in our permutation generator. Then, to handle a concrete cost function $c \in \mathcal{C}_n$, we just have to pick a random $\pi \in \Pi_n$ using $\lg_2 |\Pi_n|$ bits. This model prevents an adversary from learning Π_n and create the adversarial cost function $c^{\mathcal{U}(\Pi_n)}$ discussed

earlier. Our goal in this paper is to have a public permutation generator that works for all possible future cost functions, and this is when a linear number of random bits are needed for backwards analysis to be trusted. Such tricks can be applied whenever we have an analysis assuming full randomness, and indeed the basic message of this paper is that backwards analysis is almost as abstract as any analysis based on full randomness.

2 Preliminaries

We define $[n] = \{1, 2, \ldots, n\}$ and let S_n denote the set of permutations of $[n]$. For a graph G we let $V(G)$ and $E(G)$ denote the set of nodes and edges of G, respectively. For a set S we let $\mathcal{P}(S)$ denote the power set of S, i.e. the set containing all subsets of S.

For a random variable X and a distribution $\mathcal{D}$ we write $X \sim \mathcal{D}$ to mean that X is drawn from $\mathcal{D}$. Given X drawn from $\mathcal{D}$ the self information $\mathrm{I}(X)$ of X is defined as $\mathrm{I}(X) = -\ln\left(\Pr_{X' \sim D}(X = X')\right)$, where X and X' are i.i.d., and the entropy of X is $\mathrm{H}(X) = \mathrm{E}(\mathrm{I}(X))$. When X is drawn from $\mathcal{D}$ the entropy of $\mathcal{D}$ is defined by $\mathrm{H}(\mathcal{D}) = \mathrm{H}(X)$.

3 Transition Graphs of Distributions

Given a distribution $\mathcal{D}$ on S_n we associate a *transition graph* G defined in the following manner. The nodes of G are $\mathcal{P}([n])$, i.e., the set of all subsets of $[n]$. For every node $S \in V(G)$ the weight of S is $w_G(S) = \Pr_{\pi \sim \mathcal{D}}\left[\pi\left([|S|]\right) = S\right]$. For each non-empty set $S \subseteq [n]$ such that $w_G(S) > 0$ and $s \in S$ there is and edge from S to $S \backslash \{s\}$ with weight $w_G(S, S \backslash \{s\})$ defined by

$$w_G(S, S \backslash \{s\}) = \Pr\left[\pi\left(|S|\right) = s \mid \pi\left([|S|]\right) = S\right] .$$

Each permutation π from $\mathcal{D}$ defines, in reverse order, a walk from $[n]$ to $\varnothing$ where each edge drops an element from the current set, that is, the walk $(\pi([n]), \pi([n-1]), \ldots, \pi([1]), \varnothing)$.

When it is clear from the context which graph G we are working with, we write w instead of w_G.

Definition 3. A *transition graph* G is a graph for which there exists a distribution $\mathcal{D}$ such that G is the transition graph for $\mathcal{D}$. If G is the transition graph for $\mathcal{D}$, we write $\mathcal{D} \to G$.

The reason we introduce the transition graph is that it captures the backwards performance of the distribution. The distribution $\mathcal{D}$ is α-uniform if and only if, for each node S with positive weight, every outgoing edges $(S, S \backslash \{s\})$ has weight at most $\alpha/|S|$. Likewise, for any given cost function c, the expected cost with $\mathcal{D}$ can be computed from G as $\sum_{S, S \backslash \{s\}} (c(s, S) w(S) w(S, S \backslash \{s\}))$. The α-efficiency of $\mathcal{D}$ is thus also determined from G.

Let G_n^U be the transition graph of the uniform distribution over S_n. Because a transition graph captures backwards performance, we get that any permutation distribution $\mathcal{D}$ with transition graph G_n^U must be both backwards 1-uniform and 1-efficient. In fact, we will also argue the converse: if $\mathcal{D}$ is backwards 1-uniform or 1-efficient, then $\mathcal{D} \to G_n^U$. We will argue that this is also equivalent to $\mathcal{D}$ being minwise or maxwise, that is, we will prove Proposition 1.

Proposition 1. *Let $\mathcal{D}$ be a distribution and G_n^U the transition graph of $\mathcal{U}(S_n)$. Then the following five statements are equivalent: (i) $\mathcal{D}$ is backwards 1-uniform. (ii) $\mathcal{D}$ is (restricted) backwards 1-efficient. (iii) $\mathcal{D}$ is exact minwise. (iv) $\mathcal{D}$ is exact maxwise. (v) $\mathcal{D} \to G_n^U$.*

Proof. Let G be the distribution graph of $\mathcal{D}$. We have already seen that whether $\mathcal{D}$ is backwards 1-uniform and whether $\mathcal{D}$ is backwards 1-efficient, are properties of G. Since $\mathcal{U}(S_n)$ is backwards 1-uniform and 1-efficient this implies that (v) $\Rightarrow$ (i), (ii). Let $\pi \sim \mathcal{D}$. Then $\mathcal{D}$ is exact minwise if and only if $\Pr[\min \pi(S) = \pi(s)] = \frac{1}{|S|}$ for every $s \in S \subseteq [n]$. We can write $\Pr[\min \pi(S) = \pi(s)]$ in terms of the weights of G in the following manner

$$\Pr[\min \pi(S) = \pi(s)] = \sum_{T \subseteq [n], T \cap S = \{s\}} w_G(T) w_G(T, T \setminus \{s\}) \,. \tag{1}$$

In the same manner we note that

$$\Pr[\max \pi(S) = \pi(s)] = \sum_{T \subseteq [n], S \subseteq T} w_G(T) w_G(T, T \setminus \{s\}) \,. \tag{2}$$

Equations (1) and (2) imply that (iii) and (iv) follows from (v).

We have proved that if $G = G_n^U$ then all of (i)-(iv) hold. To prove that all the statements are equivalent we will prove that if $G \neq G_n^U$ then none of (i)-(iv) hold. So assume that G is not the distribution graph for $\mathcal{U}(S_n)$ and let $S \subseteq [n]$ and $s \in S$ be such that $w_G(S, S \setminus \{s\}) \neq w_{G_n^U}(S, S \setminus \{s\}) = \frac{1}{|S|}$. Such a S much exists since the edge weights determine the node weights. We furthermore choose S such that $|S|$ is maximal under this constraint. We also note that we may choose s such that $w_G(S, S \setminus \{s\}) > \frac{1}{|S|}$. We observe that $w_G(T) = w_{G_n^U}(T)$ for every set T with $|T| \geqslant |S|$. Here we use again, that the edge weights determine the node weights: $w_G(T)$ can be calculated by the weights of the edges on paths from $[n]$ to T, and these weights are equal in G and G_n^U.

Since $w_G(S, S \setminus \{s\}) > \frac{1}{|S|}$ it is clear that G is not backwards 1-uniform, and therefore (i) does not hold. Let $c(t, T)$ be the cost function that equals 1 when $T \neq S$, and for $T = S$ we let $c(t, T) = |S|$ if $t = s$ and $c(t, T) = 0$ for $t \neq s$. Then we have that

$$\mathrm{E}_{\pi \sim \mathcal{D}}[c(\pi)] - \mathrm{E}_{\pi \sim \mathcal{U}(S_n)}[c(\pi)] = w_G(S) w_G(S, S \setminus \{s\}) |S| - w_{G_n^U}(S) > 0 \,,$$

and therefore $\mathcal{D}$ is not backwards 1-efficient, and (ii) does not hold. Now consider the calculation of $\Pr_{\pi \sim \mathcal{D}}[\max \pi(S) = \pi(s)]$ in (2). All terms of the sum are the same if we replace G by G_n^U, except for the term $w_G(S) w_G(S, S \setminus \{s\})$. Since $w_G(S) = w_{G_n^U}(S) > 0$ and $w_G(S, S \setminus \{s\}) > w_{G_n^U}(S, S \setminus \{s\})$ we conclude that the the probability that $\max \pi(S) = \pi(s)$ is larger when $\pi \sim \mathcal{D}$ that when $\pi \sim \mathcal{U}(S_n)$ and hence $\mathcal{D}$ is non exact maxwise and (iv) does not hold.

We have thus far proved that (i), (ii), (iv), (v) are all equivalent and they imply (iii) via (v). We have proved maxwise $\Rightarrow$ minwise, so by symmetry minwise $\Rightarrow$ maxwise as well, so (iii) and (iv) are equivalent. So all of the five conditions are equivalent, which finishes the proof. □

The minimal size of a set $X \subseteq S_n$ such that $\mathcal{U}(X)$ is an exact minwise distribution was studied in [1], and by Proposition 1, this is also the bound for backwards 1-uniform and 1-effient. In [1], the minimal size was proved to be between $\mathrm{lcm}(1, 2, \ldots, n) = e^{n - o(n)}$ and 4^n. In Corollary 2 we prove that the minimal size is precisely $\mathrm{lcm}(1, 2, \ldots, n)$.

Now let G be a transition graph. We associate a *memoryless distribution* of permutations $\mathcal{D}$ to G. The permutations of $\mathcal{D}$ are constructed in the following manner. Let $[n] = A_n, A_{n-1}, \ldots, A_0$ be a random walk of length n in G starting in the node $[n]$. When the walk is at the node A the next node on the walk is chosen independently from the previous choices and randomly from the

outgoing neighbours of A, such that B is chosen with probability $w(A,B)$ if there is an edge from A to B. In this way the walk $A_n, A_{n-1}, \dots, A_0$ is a path in G such that $|A_i| = i$ for $i = 0, 1, 2, \dots, n$. Let $\pi \in S_n$ be the permutation such that $\{\pi(i)\} = A_i \backslash A_{i-1}$ for each $i \in [n]$. The distribution of π generated in the above manner is called the memoryless distribution of G.

Definition 4. Given a transition graph G the *memoryless distribution* $\mathcal{D}$ of G is the distribution on S_n described above. If $\mathcal{D}$ is the memoryless distribution of G we write $G \to \mathcal{D}$.

We remark that for a transition graph there may be multiple distributions $\mathcal{D}, \mathcal{D}'$ such that $\mathcal{D} \to G$ and $\mathcal{D}' \to G$. However, for each distribution $\mathcal{D}$ there is only one transition graph G such that $\mathcal{D} \to G$. If $G \to \mathcal{D}$ then $\mathcal{D} \to G$, but $\mathcal{D} \to G$ does not imply that $G \to \mathcal{D}$.

We also note that if $\mathcal{D} \to G \to \mathcal{D}'$ then $\mathcal{D}$ and $\mathcal{D}'$ might be different, but $G \to \mathcal{D} \to G'$ implies that $G = G'$.

We remark, without proof, that among the distributions $\mathcal{D}$ with $\mathcal{D} \to G$, the memoryless distribution of G is the one with the maximal amount of entropy.

4 Lower Bound

The main theorem of this paper, Theorem 1, shows that if a distribution $\mathcal{D}$ is α-efficient then the entropy of $\mathcal{D}$ is $\Omega\left(\frac{n}{\alpha}\right)$. The main idea is as follows. Let G be the transition graph of $\mathcal{D}$ and $\mathcal{D}'$ the memoryless distribution of G such that $\mathcal{D} \to G \to \mathcal{D}'$. We then consider the $(q+1)$-tuple $\ell(\pi) = (\pi([p]), \pi([p+1]), \dots, \pi([p+q]))$, where π is drawn from the memoryless distribution $\mathcal{D}'$, and where p, q are suitably chosen depending on $\mathcal{D}$. We then proceed to give a lower bound on $\mathrm{H}(\ell(\pi))$ in terms of n and α, and give a upper bound on $\mathrm{H}(\ell(\pi))$ in terms of $\mathrm{H}(\mathcal{D})$. The key idea is that $\mathrm{H}(\mathcal{D})$ can be used to give an upper bound on $\mathrm{H}(\ell(\pi))$ even though π is drawn from $\mathcal{D}'$ and not from $\mathcal{D}$. This comes from the fact that $\ell(\pi)$ can be described by $\pi([p])$ and $\pi([p+q])$ along with a permutation of $\pi([p+q])\backslash\pi([p])$. Each of the sets $\pi([p])$ and $\pi([p+q])$ have the same distribution whether π is chosen from $\mathcal{D}$ or $\mathcal{D}'$ since this is determined by the weights of the corresponding nodes in G, and so the entropy of each set is at most $\mathrm{H}(\mathcal{D})$. Since there are at most $q!$ permutations of q elements it follows that $\mathrm{H}(\ell(\pi)) \leqslant 2\mathrm{H}(\mathcal{D}) + \ln(q!)$. In order to give a lower bound on the entropy of $\ell(\pi)$ we use that the distribution of $\pi([p+i])$ conditioned on $(\pi([p+i+1]), \pi([p+i+2]), \dots, \pi([p+q]))$ is the same as the distribution of $\pi([p+i])$ conditioned on $\pi([p+i+1])$, because the distribution is memoryless. Informally, this means that when we discover $\pi([p+i])$ after having discovered $(\pi([p+i+1]), \pi([p+i+2]), \dots, \pi([p+q]))$ there is still a large amount of entropy in $\pi([p+i])$ - the exact amount is determined by the weights of the outgoing edges of $\pi([p+i+1])$ in G. If π had been drawn from $\mathcal{D}$ instead then this would not necessarily be the case and the proof would break down. The technical details are given below.

Theorem 1. *Let $\mathcal{D}$ be a distribution over S_n. If $\mathcal{D}$ is backwards α-efficient, then* $\mathrm{H}(\mathcal{D}) \geqslant \left\lfloor \frac{n}{48\alpha} \right\rfloor$.

We have not tried to optimize the constant in Theorem 1.

Proof. Let G be the transition graph of $\mathcal{D}$ and $\mathcal{D}'$ the memoryless distribution of G so that we have $\mathcal{D} \to G \to \mathcal{D}'$. Since $\mathcal{D}$ and $\mathcal{D}'$ have the same transition graph we conclude that $\mathcal{D}'$ is backwards α-efficient.

For each $S \subseteq [n]$ we let $v_S \in S$ be an element of S that maximizes $w(S, S\backslash\{v_S\})$. Let c be the cost function defined by

$$c(v,S) = \begin{cases} |S| & v = v_S \\ 0 & v \neq v_S \end{cases} .$$

We see that for each S we have $\mathrm{E}_{v\sim\mathcal{U}(S)}[c(v,S)] = 1$. Since $\mathcal{D}'$ is backwards α-efficient we have that $\mathrm{E}_{\pi\sim\mathcal{D}'}[c(\pi)] \leqslant \alpha n$. For each $i \in [n]$ let $t_i = \mathrm{E}_{\pi\sim\mathcal{D}'}[c(\pi(i), \pi([i]))]$, then we have $\sum_{i=1}^{n} t_i \leqslant \alpha n$.

Let $q = \left\lfloor \frac{n}{24\alpha} \right\rfloor$. If $q = 0$ then $\left\lfloor \frac{n}{48\alpha} \right\rfloor = 0$ and there is nothing to prove, so assume that $q > 0$. There are $\left\lfloor \frac{n}{2q} \right\rfloor$ ways to choose a positive integer k such that $n - kq \geqslant \frac{n}{2}$, and hence there must exist $p = n - kq \geqslant \frac{n}{2}$ such that $\sum_{i=p+1}^{p+q} t_i \leqslant \frac{1}{\left\lfloor \frac{n}{2q} \right\rfloor} \sum_{i=1}^{n} t_i \leqslant \frac{\alpha n}{\left\lfloor \frac{n}{2q} \right\rfloor} \leqslant 4\alpha q$. Fix such a p.

For a permutation π we now let $\ell(\pi)$ be the $(q+1)$-tuple defined by $\ell(\pi) = (\pi([p]), \pi([p+1]), \ldots, \pi([p+q]))$. We will bound the entropy of $\ell(\pi)$, when π is chosen from the memoryless $\mathcal{D}'$, i.e. $\mathrm{H}_{\pi\sim\mathcal{D}'}(\ell(\pi))$. The $(q+1)$-tuple $\ell(\pi)$ can be deduced from the two sets $\pi([p]), \pi([p+q])$ and the sequence $(\pi(p+1), \ldots, \pi(p+q))$. The combined entropy of the two sets is at most $\mathrm{H}_{\pi\sim D'}(\pi([p])) + \mathrm{H}_{\pi\sim D'}(\pi([p+q]))$. The sequence $(\pi(p+1), \ldots, \pi(p+q))$ is a permutation of $\pi([p+q])\backslash\pi([p])$, so given $\pi([p])$ and $\pi([p+q])$ the entropy of the permutation is no more than $\ln(q!)$. Therefore, the total entropy of $\ell(\pi)$ is at most

$$\mathrm{H}_{\pi\sim D'}(\pi([p])) + \mathrm{H}_{\pi\sim D'}(\pi([p+q])) + \ln(q!) .$$

The entropy $\mathrm{H}_{\pi\sim D'}(\pi([p]))$ can be calculated from the transition graph G of $\mathcal{D}'$. Since $\mathcal{D}$ and $\mathcal{D}'$ have the same transition graph it follows that $\mathrm{H}_{\pi\sim D'}(\pi([p])) = \mathrm{H}_{\pi\sim D}(\pi([p]))$, and the latter is bounded by $\mathrm{H}(\mathcal{D})$. Using the same argument for $\mathrm{H}_{\pi\sim D'}(\pi([p+q]))$ gives that

$$\mathrm{H}_{\pi\sim\mathcal{D}'}(\ell(\pi)) \leqslant 2\mathrm{H}(\mathcal{D}) + \ln(q!) .$$

Our goal is now to give a lower bound on the entropy of $\ell(\pi)$ when π is chosen from $\mathcal{D}'$. We note that the upper bound on the entropy of $\ell(\pi)$ holds for any distribution $\mathcal{D}_0$ such that $\mathcal{D}_0 \to G$, but for the lower bound we will use that $\mathcal{D}'$ is the memoryless distribution of G. Since $\mathcal{D}'$ is the memoryless distribution of G we have that for any $(q+1)$-tuple of sets $(X_q, X_{q+1}, \ldots, X_{q+p})$ such that $|X_i| = i$ and $X_i \subseteq X_{i+1}$ we have that

$$\begin{aligned}
&\Pr_{\pi'\sim D'}\left[\ell(\pi') = (X_q, \ldots, X_{q+p})\right] \\
&= \Pr_{\pi'\sim D'}\left[\pi'([q]) = X_q\right] \prod_{i=p+1}^{p+q} \Pr_{\pi'\sim D'}\left[\pi'([i]) = X_i \mid \pi'([i-1]) = X_{i-1}\right] \\
&= \Pr_{\pi'\sim D'}\left[\pi'([q]) = X_q\right] \prod_{i=p+1}^{p+q} w(X_i, X_{i-1}) \\
&\leqslant \prod_{i=p+1}^{p+q} w(X_i, X_{i-1}) .
\end{aligned}$$

Inserting this into the definition of the entropy of $\ell(\pi)$ gives us that

$$\mathrm{H}_{\pi\sim\mathcal{D}'}(\ell(\pi)) \geqslant -\sum_{i=p+1}^{p+q} \mathrm{E}_{\pi\sim\mathcal{D}'}\left[\ln\left(w(\pi([i]), \pi([i-1]))\right)\right] .$$

By the definition of $v_{\pi([i])}$ we have that $w(\pi([i]), \pi([i-1]))$ is bounded by $w(\pi([i]), \pi([i])\backslash v_{\pi([i])})$. Applying this together with Jensen's inequality gives us that

$$\begin{aligned}\mathrm{H}_{\pi\sim\mathcal{D}'}(\ell(\pi)) &\geqslant -\sum_{i=p+1}^{p+q} \mathrm{E}_{\pi\sim\mathcal{D}'}\left[\ln\left(w(\pi([i]), \pi([i])\backslash v_{\pi([i])})\right)\right] \\ &\geqslant -\sum_{i=p+1}^{p+q} \ln\left(\mathrm{E}_{\pi\sim\mathcal{D}'}\left[w(\pi([i]), \pi([i])\backslash v_{\pi([i])})\right]\right) .\end{aligned}$$

By the definition of c we have that

$$\begin{aligned}t_i = \mathrm{E}_{\pi\sim\mathcal{D}'}\left[c(\pi(i), \pi([i]))\right] &= i \cdot \mathrm{E}_{\pi\sim\mathcal{D}'}\left[w(\pi([i]), \pi([i])\backslash v_{\pi([i])})\right] \\ &\geqslant \frac{n}{2} \cdot \mathrm{E}_{\pi\sim\mathcal{D}'}\left[w(\pi([i]), \pi([i])\backslash v_{\pi([i])})\right] ,\end{aligned}$$

and therefore the entropy of $\ell(\pi)$ is at least

$$-\sum_{i=p+1}^{p+q} \ln\left(\frac{t_i}{n/2}\right) = q\ln\left(\frac{n}{2}\right) - \sum_{i=p+1}^{p+q} \ln(t_i) .$$

By Jensen's inequality we get that

$$\sum_{i=p+1}^{p+q} \ln(t_i) \leqslant q\ln\left(\frac{\sum_{i=p+1}^{p+q} t_i}{q}\right) \leqslant q\ln(4\alpha) .$$

Hence we have that $\mathrm{H}_{\pi\sim\mathcal{D}'}(\ell(\pi)) \geqslant q\ln\left(\frac{n}{8\alpha}\right)$. Comparing this with the upper bound on the entropy gives us that

$$2\mathrm{H}(\mathcal{D}) + \ln(q!) \geqslant q\ln\left(\frac{n}{8\alpha}\right) .$$

Using that $\ln(q!) \leqslant q\ln(q)$ and isolating $\mathrm{H}(\mathcal{D})$ gives us that

$$\mathrm{H}(\mathcal{D}) \geqslant \frac{q}{2} \cdot \ln\left(\frac{n}{8\alpha q}\right) \geqslant \frac{1}{2} \cdot \left\lfloor\frac{n}{24\alpha}\right\rfloor \cdot \ln(3) \geqslant \left\lfloor\frac{n}{48\alpha}\right\rfloor ,$$

which was what we wanted. □

5 Single Batch Lower Bound

For a set X we let $\binom{X}{k} \subseteq \mathcal{P}(X)$ be the set containing all subsets of X size k. In this section we prove following result:

Theorem 2. *Let k be an integer and $\mathcal{D}$ a distribution over $\binom{[n]}{k}$. Let $\alpha \geqslant 1$. Assume that for every cost-function c such that $\mathrm{E}_{s\sim\mathcal{U}(S')}\left[c(s, S')\right] \leqslant 1$ for any subset $S' \subseteq [n]$ of size k, it holds that:*

$$\mathrm{E}_{S\sim\mathcal{D}, s\sim\mathcal{U}([n]\backslash S)}\left[c(s, S\cup\{s\})\right] \leqslant \alpha , \tag{3}$$

If $\min\{k, n-k\} = \Omega(n)$ then either $\alpha = \Omega(n)$ or $\mathrm{H}(\mathcal{D}) = \Omega\left(\frac{n}{\alpha}\log(1+\alpha)\right)$.

Let X be a set. We say that a function $w : \mathcal{P}(X) \to [0,1]$ is *weight-function* for X. In order to prove Theorem 2 we will first prove some properties of weight-functions. We define $S(w)$ and $H(w)$ by:

$$S(w) = \sum_{A \subseteq X} w(A)\,, H(w) = \sum_{A \subseteq X} w(A) \ln\left(\frac{1}{w(A)}\right)\,.$$

In the definition of $H(w)$ we only sum over $A \subseteq X$ such that $w(A) \neq 0$. We note that if $S(w) = 1$ then w corresponds to a distribution $\mathcal{D}$ on the subsets of X. In this case $H(w)$ is the entropy of $\mathcal{D}$.

For non-negative integers i we define $\pi_i(w) = \pi_i w : \mathcal{P}(X) \to [0,1]$ by:

$$(\pi_i w)(A) = \left[w(A) \geqslant 2^{-i}\right] 2^{-i-1}\,,$$

and we let $\Delta(w) = \Delta w : \mathcal{P}(X) \to [0,1]$ be defined by $(\Delta w)(X) = 0$ and for $B \neq X$:

$$(\Delta w)(B) = \max_{a \notin B} \{w(B \cup \{a\})\}\,.$$

We note that $\Delta \pi_i w = \pi_i \Delta w$. We first prove a lemma about two basic properties of $\sum_{i \geqslant 0} S(\pi_i w)$ and $\sum_{i \geqslant 0} H(\pi_i w)$.

Lemma 1. *Let $w : \mathcal{P}(X) \to [0,1]$ be a weight-function, then it holds that:*

$$\sum_{i \geqslant 0} S(\pi_i w) \in \left[\frac{1}{2} S(w), S(w)\right] \tag{4}$$

$$\sum_{i \geqslant 0} H(\pi_i w) \leqslant 3S(w) + 2H(w)\,. \tag{5}$$

Proof. We first show (4). Let A be a subset of X with $w(A) \neq 0$, and let j be the smallest integer such that $w(A) \geqslant 2^{-j}$. Then we have that:

$$\sum_{i \geqslant 0} (\pi_i)(A) = \sum_{i \geqslant j} 2^{-j-1} = 2^{-j} \in \left[\frac{1}{2} w(A), w(A)\right]\,,$$

and the conclusion follows.

We now show (5). Let A be a subset of X with $w(A) \neq 0$ and let j be the smallest integer such that $w(A) \geqslant 2^{-j}$. Then:

$$\sum_{i \geqslant 0} (\pi_i w)(A) \ln \frac{1}{(\pi_i w)(A)} = \ln(2) \sum_{i \geqslant j} 2^{-i} \cdot i = \ln(2) \cdot 2^{-j+1} \cdot (j+1)\,.$$

We clearly have that $2^{-j+1} \geqslant w(A)$ and therefore $j - 1 \leqslant \lg \frac{1}{w(A)}$, i.e. $j + 1 \leqslant 2 + \lg \frac{1}{w(A)}$. Inserting this bound along with $2^{-j} \leqslant w(A)$ gives that:

$$\begin{aligned}
\sum_{i \geqslant 0} (\pi_i w)(A) \ln \frac{1}{(\pi_i w)(A)} &\leqslant 2w(A) \cdot \left(2\ln(2) + \ln \frac{1}{w(A)}\right) \\
&= 4\ln(2) \cdot w(A) + 2w(A) \ln \frac{1}{w(A)} \\
&< 3w(A) + 2w(A) \ln \frac{1}{w(A)}\,.
\end{aligned}$$

Using the definition of H and S gives

$$\sum_{i\geqslant 0} H(\pi_i w) = \sum_{i\geqslant 0, A\subseteq X} (\pi_i w)(A) \ln \frac{1}{(\pi_i w)(A)} \leqslant \sum_{A\subseteq X} 3w(A) + 2w(A) \lg \frac{1}{w(A)} = 3S(w) + 2H(w)\,,$$

as desired. □

For a real number x and an non-negative integer k we define $\binom{x}{k} = \frac{x(x-1)\cdots(x-k+1)}{k!}$. We will use Lovász's version of the Kruskal-Katona Theorem, see e.g. [4].

Theorem 3 (Lovász, Kruskal-Katona, [4])**.** *Let $\mathcal{A} \subseteq \binom{X}{k}$, and $|\mathcal{A}| = \binom{x}{k}$ for a real number $x \geqslant k$. Then,*

$$|\Delta\mathcal{A}| \geqslant \binom{x}{k-1}.$$

In order to make it easier to apply Theorem 3 we define $\ell_k(x)$ in the following way.

Definition 5. For an integer k and a non-negative real number x let $\ell_k(x)$ be the unique non-negative real number such that $x = \binom{k-1+\ell_k(x)}{k}$.

Using Definition 5 we can reformulate Theorem 3 in the following manner:

Corollary 1. *Let $\mathcal{A} \subseteq \binom{X}{k}$. Then,*

$$|\Delta\mathcal{A}| \geqslant |\mathcal{A}| \cdot \frac{k}{\ell_k(|\mathcal{A}|)}\,.$$

Proof. Let $x = \ell_k(|\mathcal{A}|) + k - 1$. By definition we have that $|\mathcal{A}| = \binom{x}{k}$ and so we have that:

$$|\Delta\mathcal{A}| \geqslant \binom{x}{k-1} = \binom{x}{k} \cdot \frac{k}{x-k+1} = |\mathcal{A}| \cdot \frac{k}{\ell_k(|A|)}\,,$$

as desired. □

Given $\mathcal{A} \subseteq \binom{X}{k}$ we can define a weight-function by $w_{\mathcal{A}}(A) = [A \in \mathcal{A}]$. Corollary 1 is now equivalent to $S(\Delta w_{\mathcal{A}}) \geqslant S(w_{\mathcal{A}}) \cdot \frac{k}{\ell_k(S(w_{\mathcal{A}}))}$. However, it is not clear what the correct generalisation of Corollary 1 is to weight-functions with non-uniform weights. Lemma 2 gives a bound $S(\Delta\pi_i w)$ in terms of $S(\pi_i w)$.

Lemma 2. *Let $w : \mathcal{P}(X) \to [0,1]$ be a weight-function for X with support on $\binom{X}{k}$. If $S(w) \leqslant 1$ and $w(A) = 0$, then for every non-negative integer i:*

$$S(\Delta\pi_i w) \geqslant S(\pi_i w) \cdot \frac{k}{\ell_r(2^i)}\,.$$

Proof. Let $\mathcal{A}$ denote the support of $\pi_i w$. Clearly $|\mathcal{A}| \leqslant 2^i S(w) \leqslant 2^i$ by the definition of π_i. We note that $(\Delta\pi_i w)(B) = 2^{-i-1} \cdot [B \in \Delta\mathcal{A}]$. By Theorem 3 we have that $|\Delta\mathcal{A}| \geqslant |\mathcal{A}| \cdot \frac{k}{\ell_k(|\mathcal{A}|)}$. We clearly have that

$$S(\pi_i w) = |\mathcal{A}| \cdot 2^{-i-1},\ S(\Delta\pi_i w) = |\Delta\mathcal{A}| \cdot 2^{-i-1}\,.$$

Inserting this gives

$$S(\Delta\pi_i w) \geqslant S(\pi_i w) \cdot \frac{k}{\ell_k(2^i)}\,,$$

as desired. □

Lemma 3. *Let $w : \mathcal{P}(X) \to [0,1]$ be a weight-function with $S(w) = 1$ and let k be a non-negative integer the support of w is contained in $\binom{X}{k}$. Assume that $S(\Delta w) \leqslant \alpha$. Then either $\alpha = \Omega(k)$ or $H(w) = \Omega\left(\frac{k}{\alpha}\log(1+\alpha)\right)$.*

Proof. Since $S(\Delta w) \leqslant \alpha$ we conclude by Lemma 1 that

$$\sum_{i\geqslant 0} S(\Delta\pi_i w) = \sum_{i\geqslant 0} S(\pi_i \Delta w) \leqslant S(\Delta w) \leqslant \alpha\,.$$

We also see that $\sum_{i\geqslant 0} S(\pi_i w) \geqslant \frac{1}{2}S(w) = \frac{1}{2}$. Therefore there exists a largest integer j such that $\sum_{i\geqslant j} S(\pi_i w) \geqslant \frac{1}{4}$. Since j is the largest such integer we also have that $\sum_{i=0}^{j} S(\pi_i w) \geqslant \frac{1}{4}$. Note that by Lemma 1:

$$2H(w) + 3 \geqslant \sum_{i\geqslant 0} H(\pi_i w) = \sum_{i\geqslant 0} S(\pi_i w)(i+1) \geqslant (j+1)\sum_{i\geqslant j} S(\pi_i w) \geqslant \frac{1}{4}(j+1)\,,$$

and therefore we either have $j = O(1)$ or $H(w) = \Omega(j)$.

On the other hand Lemma 2 gives that

$$\alpha \geqslant \sum_{i=0}^{j} S(\Delta\pi_i w) \geqslant \sum_{i=0}^{j} S(\pi_i w)\frac{k}{\ell_k(2^i)} \geqslant \frac{1}{4}\cdot\frac{k}{\ell_k(2^j)}\,,$$

i.e. $\ell_k(2^j) \geqslant \frac{k}{4\alpha}$. Let $\beta = \left\lfloor \frac{k}{4\alpha} \right\rfloor - 1$. If $\beta < \frac{k}{8\alpha}$ or $j = O(1)$ then we have $\alpha = \Omega(k)$ as desired. So assume that $\beta \geqslant \frac{k}{8\alpha}$ and $H(w) = \Omega(j)$. By definition of $\ell_k(2^j)$ we have that:

$$2^j = \binom{k-1+\ell_k(2^j)}{k} \geqslant \binom{k+\beta}{k} = \binom{k+\beta}{\beta} \geqslant \left(\frac{k+\beta}{\beta}\right)^{\beta}\,.$$

Taking the logarithm on both sides yields:

$$j \geqslant \lg\left(1+\frac{k}{\beta}\right)\cdot\beta \geqslant \frac{k}{8\alpha}\cdot\lg(1+4\alpha)\,.$$

Since $H(w) = \Omega(j)$ we get $H(w) = \Omega\left(\frac{k}{\alpha}\log(1+\alpha)\right)$ which was what we wanted. □

We now use Lemma 3 to prove Theorem 2.

Proof of Theorem 2. Let $w : \mathcal{P}([n]) \to [0,1]$ be defined by $w(A) = \Pr_{S\sim\mathcal{D}}[S = [n]\backslash A]$. Then w is a weigh-function with $S(w) = 1$ and $H(w) = \mathrm{H}(\mathcal{D})$. The support of w is contained in $\binom{[n]}{n-k}$.

Let $f : \mathcal{P}(X) \to X$ be a mapping such that for every $B \in \binom{[n]}{k+1}$ we have $f(B) \in B$ and

$$w\left(B\backslash\{f(B)\}\right) = \max_{b\in B}\{w\left(B\backslash\{b\}\right)\}\,,$$

that is, $b = f(B)$ maximizes $w(B \setminus \{b\})$ constrained to $b \in B$. We define the cost function $c(b, B)$ for all $B \in \binom{[n]}{k+1}$ and $b \in B$ by $c(b, B) = k + 1$ if $b = f(B)$ and $c(b, B) = 0$ otherwise. By definition (3) must hold for this cost function. We plug c into (3) to obtain

$$\frac{1}{n-k} \sum_{A \in \binom{[n]}{k}} \Pr_{S \sim \mathcal{D}}[S = A] \sum_{a \notin A} c(a, A \cup \{a\}) \leqslant \mathrm{E}_{S \sim \mathcal{D}, s \sim \mathcal{U}([n] \setminus S)}[c(s, S \cup \{s\})] \leqslant \alpha \,.$$

We note that by the definition of $f(B)$ we get:

$$\begin{aligned} \frac{1}{n-k} \sum_{A \in \binom{[n]}{k}} \Pr_{S \sim \mathcal{D}}[S = A] \sum_{a \notin A} c(a, A \cup \{a\}) &= \frac{1}{n-k} \sum_{B \in \binom{[n]}{k+1}} \sum_{b \in B} \Pr_{S \sim \mathcal{D}}[S = B \setminus \{b\}] \, c(b, B) \\ &= \frac{k+1}{n-k} \sum_{B \in \binom{[n]}{k+1}} \Pr_{S \sim \mathcal{D}}[S = B \setminus \{f(B)\}] \end{aligned}$$

By the definition of Δw we get that:

$$\Pr_{S \sim \mathcal{D}}[S = B \setminus \{f(B)\}] = \max_{b \in B} \left\{ \Pr_{S \sim \mathcal{D}}[S = B \setminus \{b\}] \right\} = \max_{b \notin [n] \setminus B} \{w(([n] \setminus B) \cup \{b\})\} = (\Delta w)([n] \setminus B) \,.$$

That is,

$$\alpha \geqslant \frac{k+1}{n-k} \sum_{B \in \binom{[n]}{k+1}} \Pr_{S \sim \mathcal{D}}[S = B \setminus \{f(B)\}] = \frac{k+1}{n-k} \sum_{B' \in \binom{[n]}{n-k-1}} (\Delta w)(B') = \frac{k+1}{n-k} S(\Delta w) \,.$$

Since $k = \Omega(n)$ we get that $S(\Delta w) = O(\alpha)$. Now Lemma 3 gives that either $\alpha = \Omega(n-k)$ or $\mathrm{H}(\mathcal{D}) = H(w) = \Omega\left(\frac{n-k}{\alpha} \log(1+\alpha)\right)$. Since $n - k = \Omega(n)$ this is exactly what we wanted to prove. □

6 Upper bound

We present a general lemma about transition graphs that will help us find distributions that are uniform on relatively small sets that have desirable properties. The lemma itself is fairly straightforward and similar statements might be known.

Lemma 4. *Let G be a transition graph, and let δ_V and δ_E be the minimal non-zero weight of a node and an edge in G, respectively, and assume that $w_G(S, S')^{-1}$ is an integer for every edge (S, S') with non-zero weight. Let p be the probability of the most probable permutation in the memoryless distribution of G. Let $\varepsilon \in (0, 1]$, and t be integer satisfying $\frac{1}{p} \geqslant t \geqslant \frac{8n}{\delta_V \delta_E \varepsilon}$. There exists a set X such that the transition graph G' of $\mathcal{U}(X)$ satisfies: For each node S, $w_{G'}(S) = (1 \pm \varepsilon) w_G(S)$ and for each edge (S, S'), $w_{G'}(S, S') = \left(1 \pm \frac{\varepsilon}{4n}\right) w_G(S, S')$.*

Furthermore, if $t \cdot w_G(S) \cdot w_G(S, S')$ is an integer for each edge (S, S') of G then $G' = G$.

Proof. Put t pebbles on the node $[n]$ of G and 0 pebbles on each other node of G. We now run the following algorithm: For each $i = n, n-1, n-2, \ldots, 1$ do the following. For each subset $S \subseteq [n]$ of size i consider the node S of G, and say that there are currently s pebbles on S. Now let $T_1, \ldots, T_i$ be all subsets of S of size $i - 1$. We then move all the pebbles from S to the nodes $T_1, \ldots, T_i$ such

that we move $\approx sw_G(S, T_i)$ pebbles from S to T_i. That is, we move at least $\lfloor sw_G(S, T_i) \rfloor$ pebbles and at most $\lceil sw_G(S, T_i) \rceil$ pebbles from S to T_i. We do this in the following way. Let V be the set of pebbles in S and partition V as $V = V_1 \cup \dots V_k$ where each V_i holds a set of pebbles that have all followed the same path from $[n]$ to S. Then we use Lemma 8 with $\delta_i = w_G(S, T_i)$ to create a function $f : V \to [k]$. A pebble $v \in V$ is then moved to $T_{f(v)}$.

When the algorithm is finished all the pebbles are at the node $\varnothing$. Each pebble went from $[n]$ to $\varnothing$ through a path $[n] = S_n, S_{n-1}, \dots, S_0 = \varnothing$. The pebble then correspond to the unique permutation satisfying that $\{\pi(i)\} = S_i \backslash S_{i-1}$. For any permutation π, since we used Lemma 8 to move the pebbles, the number of pebbles that took the path corresponding to π, say $S_n, S_{n-1}, \dots, S_0$, is at most

$$\lceil w_G(S_1, S_0) \lceil \dots \lceil w_G(S_n, S_{n-1}) t \rceil \dots \rceil \rceil = \lceil w_G(S_1, S_0) \cdots w_G(S_n, S_{n-1}) \cdot t \rceil ,$$

which equals $\lceil p't \rceil$, where p' is the probability of getting π when drawing a permutation from the memoryless distribution of G. Since $p't \leqslant 1$ by assumption this shows that any two pebbles followed different paths in G and therefore correspond to different permutations. Let $\mathcal{D}$ be the uniform distribution of the permutations corresponding to the pebbles, and let G' be the transition graph of $\mathcal{D}$.

We will prove that $w_{G'}(S, S') = \left(1 \pm \frac{\varepsilon}{4n}\right) w_G(S, S')$ for all edges (S, S'). Assume for the sake of contradiction that this is false, and let (S, S') be an edge for which it does not hold with the largest possible value of $|S|$. We can calculate the weight of S in G' by considering all paths from $[n]$ to S and summing the products of the weights of the edges on each path. These edge weights are at most a factor $1 - \frac{\varepsilon}{4n}$ smaller than the corresponding edge weights in G, which can be used to calculate $w_G(S)$. This implies that the weight of S in G' is at least $\left(1 - \frac{\varepsilon}{4n}\right)^{n-|S|} w_G(S) \geqslant \left(1 - \frac{\varepsilon}{4n}\right)^{n} w_G(S)$, and hence $w_{G'}(S)$ is at least $(1 - \varepsilon/2)\, w_G(S) \geqslant \frac{1}{2} w_G(S) \geqslant \frac{1}{2}\delta_V$. The number of pebbles that where in S was exactly $tw_{G'}(S)$. So the number of pebbles that where moved from S to S' is either $\lfloor tw_{G'}(S) w_G(S, S') \rfloor$ or $\lceil tw_{G'}(S) w_G(S, S') \rceil$. The weight $w_{G'}(S, S')$ is therefore at most

$$\begin{aligned} \frac{\lceil tw_{G'}(S) w_G(S, S') \rceil}{tw_{G'}(S)} < \frac{1 + tw_{G'}(S) w_G(S, S')}{tw_{G'}(S)} &= w_G(S, S') + \frac{1}{tw_{G'}(S)} \\ &\leqslant \left(1 + \frac{\varepsilon}{4n}\right) w_G(S, S') , \end{aligned}$$

where the last inequality follows from the fact that $\frac{1}{tw_{G'}(S)} \leqslant \frac{2}{tw_G(S)}$ and the definition of t. We get that $w_{G'}(S, S')$ is at least $\left(1 - \frac{\varepsilon}{4n}\right) w_G(S, S')$ in the same manner. A contradiction. So the assumption was wrong and $w_{G'}(S, S') = \left(1 \pm \frac{\varepsilon}{4n}\right) w_G(S, S')$ holds for every edge of G'.

By using the fact that we can calculate the node weights from the edge weights it is apparent that $w_{G'}(S) = \left(1 \pm \frac{\varepsilon}{4n}\right)^n w_G(S)$ which clearly implies $w_{G'}(S) = (1 \pm \varepsilon)\, w_G(S)$ as desired. □

Let G be the transition graph of $\mathcal{U}(S_n)$ and let $t = \operatorname{lcm}_{k \in [n]} \left\{ \binom{n}{k} \cdot k \right\}$. For each edge (S, S') of G we have $w_G(S) w_G(S, S') = \binom{n}{k}^{-1} \cdot k^{-1}$ and $tw_G(S) w_G(S, S')$ is an integer. Therefore $G' = G$ in Lemma 4, and there exists a distribution, which is uniform on a set of t permutations that has the same transition graph as $\mathcal{U}(S_n)$. By Proposition 1 this distribution is exact minwise and exact backwards uniform as well. It is an easy exercise to prove that $t = \operatorname{lcm}(1, 2, \dots, n)$, and the proof is deferred to Lemma 7 The result is summarized below.

Corollary 2. *There exists a subset $X \subseteq S_n$ of size $\operatorname{lcm}(1, 2, \dots, n) = e^{n - o(n)}$ such that $\mathcal{U}(X)$ is exact minwise and exact backwards uniform.*

This is tight, since it is shown in [1] that for any uniform distribution on X that is exact minwise, the size of X is divisible by $\operatorname{lcm}(1, 2, \ldots, n)$. This also improves on the previous upper bound from [1] of $4^{n-O(\log^2 n)}$.

Now we will show that for any $\alpha > 1$ there exists a backwards α-uniform distribution $\mathcal{D}$ that is the uniform distribution on $\exp\left(O\left(\frac{n}{\alpha} \cdot \left(1 + \log^2 \alpha\right)\right)\right)$ permutations. We will do this by first describing a preliminary distribution $\mathcal{D}'$, and then convert $\mathcal{D}'$ into a similar distribution $\mathcal{D}$ using Lemma 4.

For simplicity we assume that n and α are powers of two. Let $t = \frac{n}{\alpha}$. We now describe an algorithm that takes a positive integer k and a set S of $2^k t$ integers and outputs a permutation of S. We describe the permutation, that the algorithm finds when $S = [2^k t]$. For general S we order the elements of S as $S = \{s_1, s_2, \ldots, s_{2^k t}\}$ such that $s_1 < s_2 < \ldots < s_{2^k t}$. Then we find a random permutation π' of $[2^k t]$, and the algorithm outputs $\pi : S \to S$ defined by $\pi(s_i) = s_{\pi(i)}$. So from now on we may assume that $S = [2^k t]$. If $k = 1$ the algorithm simply chooses a permutation uniformly at random S and outputs that permutation. If $k > 1$ the algorithm partitions S into two sets $S_0 = [2^{k-1} t]$ and $S_1 = [2^k t] \backslash [2^{k-1} t]$ consisting of the smallest and largest half of S, respectively. The algorithm now recursively finds permutations of S_0 and S_1, which we call π_0 and π_1, respectively. We assume that π_0 and π_1 are independent. Let S' be the set of the t last elements of π_0 and π_1, i.e.,

$$S' = \pi_0\left(S_0 \backslash [2^{k-1} t - t]\right) \cup \pi_1\left(S_1 \backslash [2^t - t]\right) .$$

Let $s'_1, \ldots, s'_{2t}$ be a uniformly random ordering of S'. The result of the algorithm is now obtained by concatenating the first $2^{k-1} t - t$ elements of π_0 with the first $2^{k-1} t - t$ elements of π_1 and then appending $s'_1, \ldots, s'_{2t}$. That is, the output of the algorithm is the following permutation:

$$\pi(i) = \begin{cases} \pi_0(i) & i \leqslant 2^{k-1} t - t \\ \pi_1\left(i - (2^{k-1} t - t)\right) & 2^{k-1} t - t < i \leqslant 2^k t - 2t \\ s'_{i-(2^k - 2t)} & 2^k - 2t < i \end{cases} .$$

For each value of k this gives a distribution of permutations on $\left\{1, 2, \ldots, 2^k t\right\}$ which we call $\mathcal{D}_k$.

We now prove Lemma 5 and Lemma 6 which give a lower bound on the non-zero weights of the nodes in the transition graph of $\mathcal{D}_k$.

Lemma 5. *Let k be a positive integer and let $X \subseteq [2^k t]$ be a set of size s. The probability that all elements of X are among the last t elements in a permutation chosen randomly according to $\mathcal{D}_k$ is at least 2^{-2sk}.*

Proof. We prove the claim by induction on k. First say that $k = 1$. Then the probability that all elements from X are among the last t elements is exactly $\frac{\binom{t}{s}}{\binom{2t}{s}} > 2^{-2s}$ as desired.

Now say that $k > 1$, and assume that the claim holds for the distribution $\mathcal{D}_{k-1}$. Let $A_0 = [2^{k-1} t], A_1 = [2^k t] \backslash [2^{k-1} t]$ and set $X_0 = X \cap A_0, X_1 = X \cap A_0$. The algorithm first computes two permutations of A_0 and A_1. The probability that all elements of X_i are among the last t elements in the permutation of A_i is at least $2^{-2(k-1)|X_i|}$ by the induction hypothesis. So by the independence of the two permutations, the probability that all elements of X_i are among the last t elements in the permutation of A_i for $i = 0$ and $i = 1$ is at least $2^{-2(k-1)|X_0|} \cdot 2^{-2(k-1)|X_1|} = 2^{-2(k-1)s}$. That is, the probability that all elements of X are among the last $2t$ elements of the permutation is at

least $2^{-2(k-1)s}$. Now assume that this is the case. Then using the same analysis as in the case $k = 1$ we see that that the probability that all elements of X are among the last t elements in the permutation is at least $\frac{\binom{t}{s}}{\binom{2t}{s}} > 2^{-2s}$. So the probability that all elements of X are contained among the last t elements in the permutation must be at least $2^{-2(k-1)s} \cdot 2^{-2s} = 2^{-2ks}$ as desired. □

Lemma 6. *Let k be a positive integer and $X \subseteq [2^k t]$ a set of s positive integers. The probability the the set of the first s elements is X in a permutation chosen randomly according to $\mathcal{D}_k$ is either 0 or at least $2^{-2k^2 t}$.*

Proof. We prove the claim by induction on k. First assume that $k = 1$. Then the probability is exactly $\frac{1}{\binom{2t}{s}} > 2^{-2t} = 2^{-2k^2 t}$ as desired.

Now say that $k > 1$ and assume that the claim holds for $\mathcal{D}_{k-1}$. Let $A_0 = [2^{k-1}t], A_1 = [2^k t] \backslash [2^{k-1}t]$ and set $X_0 = X \cap A_0, X_1 = X \cap A_0$. The algorithm first recursively computes permutations π_i of A_i for $i = 0, 1$ and then uses them to compute a permutation π of $[2^k t]$. We now consider three cases:

(1) $s \leqslant (2^{k-1} - 1)t$.

(2) $(2^{k-1} - 1)t < s \leqslant (2^k - 2)t$.

(3) $(2^k - 2)t < s$.

Case (1): Whether the first s elements of π is X depends only on π_0, and by the induction hypothesis the probability is either 0 or at least $2^{-2(k-1)^2 t} > 2^{-2k^2 t}$.

Case (2): The first $(2^{k-1} - 1)t$ elements of π are elements from A_0 and the next $s - (2^{k-1} - 1)t$ elements are elements from A_1. This means that if $|X_0| \neq (2^{k-1} - 1)t$ the probability is 0, so assume that $|X_0| = (2^{k-1} - 1)t$. Now the first s elements of π is X if and only if the last t elements of the π_0 is $A_0 \backslash X_0$ and the first $s - (2^{k-1} - 1)t$ elements of the permutation of A_1 is X_1. By Lemma 5 and the induction hypothesis applied to $\mathcal{D}_{k-1}$ this happens with probability at least $2^{-2(k-1)t} \cdot 2^{-2(k-1)^2 t} > 2^{-2k^2 t}$.

Case (3): The set of the first s elements of π is X exactly if the last elements of π are $[2^k t] \backslash X = A_0 \backslash X_0 \cup A_1 \backslash X_1$. If $A_i \backslash X_i$ contains more than t elements for $i = 0$ or $i = 1$ the probability that this happens is 0, so assume that $|A_i \backslash X_i| \leqslant t$. The last elements of π are $A_0 \backslash X_0 \cup A_1 \backslash X_1$ only if the last t elements of π_i contains $A_i \backslash X_i$ for $i = 0, 1$. The latter happens with probability at least $\left(2^{-2(k-1)t}\right)^2$ by Lemma 5. If this happens then the probability that the last of π is $A_0 \backslash X_0 \cup A_1 \backslash X_1$ $\binom{2t}{|A_0 \backslash X_0 \cup A_1 \backslash X_1|}^{-1} > 2^{-2t}$. So the probability that the first s elements of π is X is either 0 or at least $\left(2^{-2(k-1)t}\right)^2 \cdot 2^{-2t} > 2^{-2k^2 t}$.

So in each case the probability is either 0 or at least $2^{-2k^2 t}$. By induction this holds for all positive integers k. □

Lemma 6 shows that the weight of any node in the transition graph of $\mathcal{D}_{\lg \alpha}$ is either 0 or at least $\exp\left(-O\left(\frac{n}{\alpha} \cdot \left(1 + \log^2 \alpha\right)\right)\right)$. Combining this with Lemma 4 shows that there exists a set X of size $\exp\left(O\left(\frac{n}{\alpha} \cdot \left(1 + \log^2 \alpha\right)\right)\right)$ such that $\mathcal{U}(X)$ is backwards α-uniform.

References

[1] Andrei Z. Broder, Moses Charikar, Alan M. Frieze, and Michael Mitzenmacher. Min-wise independent permutations. *Journal of Computer and System Sciences*, 60(3):630–659, 2000. Announced at STOC'98.

[2] Timothy M Chan. Backwards analysis of the karger-klein-tarjan algorithm for minimum spanning trees. *Information Processing Letters*, 67(6):303–304, 1998.

[3] Mark de Berg, Otfried Schwarzkopf, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2nd edition, 2000.

[4] Péter Frankl. A new short proof for the kruskal-katona theorem. *Discrete Mathematics*, 48(2-3):327–329, 1984.

[5] C. A. R. Hoare. Quicksort. *Comput. J.*, 5(1):10–15, 1962.

[6] Piotr Indyk. A small approximately min-wise independent family of hash functions. *Journal of Algorithms*, 38(1):84–90, 2001. Announced at SODA'99.

[7] David R. Karger, Philip N. Klein, and Robert Endre Tarjan. A randomized linear-time algorithm to find minimum spanning trees. *J. ACM*, 42(2):321–328, 1995.

[8] Rajeev Motwani and Prabhakar Raghavan. *Randomized algorithms*. Cambridge University Press, 1995.

[9] Ketan Mulmuley. Randomized geometric algorithms and pseudorandom generators. *Algorithmica*, 16(4/5):450–463, 1996.

[10] Seth Pettie and Vijaya Ramachandran. Randomized minimum spanning tree algorithms using exponentially fewer random bits. *ACM Trans. Algorithms*, 4(1):5:1–5:27, 2008.

[11] Michael O. Rabin. Probabilistic algorithms. In J. F. Traub, editor, *Algorithms and Complexity*, pages 21–30. Academic Press, 1976.

[12] Raimund Seidel. Backwards analysis of randomized geometric algorithms. In János Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, chapter 2, pages 37–67. Springer-Verlag, 1993.

[13] Mikkel Thorup. Fast and powerful hashing using tabulation. *CoRR*, abs/1505.01523, 2015.

[14] Mark N. Wegman and Larry Carter. New classes and applications of hash functions. *Journal of Computer and System Sciences*, 22(3):265–279, 1981. Announced at FOCS'79.

7 Appendix

Lemma 7. *For any integer n we have that:*

$$\operatorname{lcm}_{k\in[n]}\left(\binom{n}{k}\cdot k\right) = \operatorname{lcm}(1,2,\ldots,n)\,. \tag{6}$$

Proof. Fix n and let $a = \operatorname{lcm}_{k\in[n]}\left(\binom{n}{k}k\right)$ and $b = \operatorname{lcm}(1,2,\ldots,n)$. For a prime p and a integer m we let $v_p(m)$ denote the largest integer ℓ such p^ℓ divides m. It is well-known that $v_p(m!)$ can be calculated by the following sum:

$$v_p(m!) = \sum_{i\geqslant 1}\left\lfloor \frac{m}{p^i}\right\rfloor\,. \tag{7}$$

Using that $\binom{n}{k} = \frac{n!}{(n-k)!(k-1)!}$ and (7) we get that

$$v_p(a) = \max_{k\in[n]}\left\{\sum_{i\geqslant 1}\left(\left\lfloor\frac{n}{p^i}\right\rfloor - \left\lfloor\frac{n-k}{p^i}\right\rfloor - \left\lfloor\frac{k-1}{p^i}\right\rfloor\right)\right\}\,. \tag{8}$$

We will prove that $v_p(a) = v_p(b)$ for every prime p. This clearly implies that $a = b$. Fix a prime p, and let $\ell = v_p(b)$. Then $p^\ell \leqslant n$ and $p^{\ell+1} \geqslant n$. Therefore we can write $v_p(a)$ as:

$$\begin{aligned} v_p(a) &= \max_{k\in[n]}\left\{\sum_{i=1}^{\ell}\left(\left\lfloor\frac{n}{p^i}\right\rfloor - \left\lfloor\frac{n-k}{p^i}\right\rfloor - \left\lfloor\frac{k-1}{p^i}\right\rfloor\right)\right\} \\ &\leqslant \max_{k\in[n]}\left\{\sum_{i=1}^{\ell}\left(\left(\frac{n}{p^i}+\frac{p^i-1}{p^i}\right) - \frac{n-k}{p^i} - \frac{k-1}{p^i}\right)\right\} = \ell\,. \end{aligned}$$

To prove that $v_p(a) \geqslant \ell$ we let $k = p^\ell$ in (8) to obtain that

$$\begin{aligned} v_p(a) &\geqslant \sum_{i=1}^{\ell}\left(\left\lfloor\frac{n}{p^i}\right\rfloor - \left\lfloor\frac{n-p^\ell}{p^i}\right\rfloor - \left\lfloor\frac{p^\ell-1}{p^i}\right\rfloor\right) \\ &= \sum_{i=1}^{\ell}\left(\left\lfloor\frac{n}{p^i}\right\rfloor - \left(\left\lfloor\frac{n}{p^i}\right\rfloor - p^{\ell-i}\right) - \left(p^{\ell-i}-1\right)\right) = \ell\,. \end{aligned}$$

Hence $v_p(a) = \ell = v_p(b)$, and since this holds for every prime p we have $a = b$ and (6) holds as desired. □

Lemma 8. *Let $V = V_1 \cup V_2 \cup \ldots V_k$ be a partition of a set V. Let r be an integer and let $\delta_1, \delta_2, \ldots, \delta_r \geqslant 0$ be non-negative real numbers such that $\sum_{i=1}^{r}\delta_i = 1$. Then there exists a function $g : V \to [r]$ such that for all $i \in [r]$ and $j \in [k]$ we have*

$$\lfloor \delta_i\,|V_j|\rfloor \leqslant \left|g^{-1}(i)\cap V_j\right| \leqslant \lceil \delta_i\,|V_j|\rceil\,,$$

and such that it also holds that

$$\lfloor \delta_i\,|V|\rfloor \leqslant \left|g^{-1}(i)\cap V\right| \leqslant \lceil \delta_i\,|V|\rceil$$

Proof. Let G be a with the nodes $U = \{s, t\} \cup \{V_1, \ldots, V_k\} \cup [r]$, where s and t are a source and a sink, respectively. For each $j \in [k]$ let $\alpha_j = |V_j| - \sum_{i \in [r]} \lfloor \delta_i |V_j| \rfloor$, and let $\beta_i = \lceil |V| \delta_i \rceil - \sum_{j \in [k]} \lfloor \delta_i |V_j| \rfloor$. The graph has an edge from s to V_j with capacity α_j for every $j \in [k]$, and an edge from i to t with capacity β_i for every $i \in [r]$. For each pair (V_j, i) such that $\delta_i |V_j|$ is not an integer we add an edge from V_j to i with capacity 1.

We will now show that the maximum st-flow in G is $\alpha = \sum_{j \in [k]} \alpha_j$. Since the cut $(s, U \setminus \{s\})$ has value α, the flow is at most α. On the other hand let f be a flow function defined as follows. For any $j \in [k]$ $f(s, V_j) = \alpha_j$. For any $j \in [k], i \in [r]$ $f(V_j, i) = \delta_i |V_j| - \lfloor \delta_i |V_j| \rfloor$. For any $i \in [r]$ $f(i, t) = |V| \delta_i - \sum_{j \in [k]} \lfloor \delta_i |V_j| \rfloor$. It is easy to verify that f is a flow with value α.

Since all the capacities of G are integers, there exists a maximum st-flow f_I where the flow along each edge is an integer. We now construct g from f_I. For each $j \in [k]$ we partition V_j into $V_j = V_j^1 \cup V_j^2 \cup \ldots \cup V_j^r$ such that $\left|V_j^i\right| = \lfloor \delta_i |V_j| \rfloor + f_I(V_j, i)$ for every $i \in [r]$. This is clearly possible. Now we define g such that $g(v) = i$ for every $v \in \bigcup_{j \in [k]} V_j^i$, and it is easy to see that g. □